# Iterative Partial Rounding for Vertex Cover with Hard Capacities

Mong-Jen Kao

Institute of Information Science,
Academia Sinica, Taiwan.
mong@iis.sinica.edu.tw

arXiv:1606.08667v4 [cs.DS] 14 Oct 2016

**Abstract**

We provide a simple and novel algorithmic design technique, for which we call *iterative partial rounding*, that gives a tight rounding-based approximation for vertex cover with hard capacities (VC-HC). In particular, we obtain an $f$-approximation for VC-HC on hypergraphs, improving over a previous results of Cheung et al. (SODA 2014) to the tight extent. This also closes the gap of approximation since it was posted by Chuzhoy and Naor in (FOCS 2002). We believe that our rounding technique is of independent interest when hard constraints are considered.

Our main technical tool for establishing the approximation guarantee is a separation lemma that certifies the existence of a strong partition for solutions that are basic feasible in an extended version of the natural LP.

**Keywords:** iterative partial rounding, capacitated vertex cover, hard capacities, approximation algorithm

# 1 Introduction

We consider the vertex cover problem with hard capacity constraints (VC-HC) on hypergraphs. In this problem, we are given a hypergraph $G = (V, E)$ with a maximum edge size $f$. Each $e \in E$ is associated with an edge demand $d_e$ and each $v \in V$ is associated with a capacity $c_v$ and an available multiplicity (the number of available copies) $m_v$. The objective is to find a minimum-size vertex cover, a vertex multi-set represented by an assignment function, such that the demands of the edges can be covered by the capacities of the vertices chosen in the multiset and the multiplicity of each vertex does not exceed its available multiplicity. When no upper bound is imposed on the multiplicities of each vertex, i.e., $m(v) = \infty$ for all $v$, this problem is then referred to as (soft) capacitated vertex cover (CVC). In this case weighted vertex set is usually considered and minimum-weight vertex multi-set is sought. In this paper, we assume that VC-HC takes unweighted vertices unless specified otherwise.

The study of VC-HC was initiated in the notable work of Chuzhoy and Naor [7], where normal graphs with unit edge demand are assumed, i.e., $f = 2$ and $d_e = 1$ for all $e$. Under this setting, they established a surprising result that, while this setting admits constant factor approximations, it becomes set-cover hard when $\{0, 1\}$-weighted vertices are considered, i.e., the weight of each vertex can be either 0 or 1. This implies an interesting logarithmic separation on the approximability between weighted and unweighted vertices. In the same work, the status of (unweighted) VC-HC with general edge demand was left as an open problem. While the gap of approximation for this problem was settled partially a decade later by Saha and Khuller [28] and Cheung et al [6], the exact approximability of this problem remains unsettled.

For a brief introduction on the current status of VC-HC, Chuzhoy and Naor [7] presented a 3-approximation for normal graphs with unit edge demand. This result was later improved by Gandhi et al [11] to a tight 2-approximation with a refined approach. Saha and Khuller [28] considered general edge demands and presented an $O(f)$-approximation for hypergraphs. This result was improved by Cheung et al [6] with a $\left(1 + 2/\sqrt{3}\right) \approx 2.155$-approximation for normal graphs and a $2f$-approximation for hypergraphs.

One intriguing thing in the development of VC-HC is how the techniques that were used to solve this problem are influenced (constrained) by the complexity separation between weighted and unweighted vertex sets. Therefore, if an approach were to work, it has to be sensitive enough to tell the difference between the assumptions. This nature, as also pointed out in [6], renders existing techniques for CVC, such as primal-dual [15], dependent rounding [13], LP rounding, etc., not directly applicable to VC-HC since very often they are not sensitive to the weight of the vertices.

In fact, all of existing results for VC-HC are built on the same two-staged rounding principle: First, it begins with a vertex-side threshold rounding. Then one or multiple edge-vertex patching procedure is introduced to meet the covering guarantee. The main challenge of this approach has been on devising a delicate patching procedure. This has been demonstrated in the the works [6, 28]. In particular, a near-tight 2.155-approximation for normal graphs is obtained, using an elegant matching structure extracted during the patching stage, combined with a neat interaction back to the threshold parameter of the first stage. Although improved approximations are obtained, it seems that current two-staged rounding techniques have reached their limitations, and significant new ideas are required to close the gap.

Therefore, a natural and central question that arises is thus:

*Can the rounding for VC-HC be done without patching?*

In this work, we provide a positive answer to the above question: We present a novel rounding-based approximation algorithm for VC-HC which closes the gap of approximation since it was posted a decade ago. Compared to prior results, our algorithm is very simple to state and easy

to understand. We believe that our rounding technique is also of independent interest when hard constraints are considered.

### 1.1 Our Results

Our approximation algorithm is stated as follows.

**Theorem 1.** There is a rounding-based approximation algorithm for VC-HC that, given an instance with largest size of hyperedges $f \geq 2$, produces an $f$-approximation in polynomial-time.

Our algorithm is driven by a simple *iterative partial rounding* scheme. In each iteration, it makes partial decisions based on current working LP and rounds the vertices *fractionally*. When no such decisions are there to be made, it rounds up all vertices *unconditionally* and stops.

During this process, the input instance along with the working LP are modified gradually. While the edges may be removed (folded) from the instance as the algorithm iterates, the lower bounds on the multiplicity of the vertices are meanwhile strengthened. The intermediate solutions to the working LPs may not be feasible for the original input instance. However, we guarantee that the overall optimality is not lost before the final rounding is done, and all the partial decisions we made entail no potential loss in the feasibility and approximation guarantee of the final solution. Together this gives our tight approximation for this problem.

In contrast to previous two-staged approaches, instead of making one big decisive rounding followed with patching, our algorithm makes only partial indecisive moves which are also proven harmless and waits before the final decisive rounding can be made. From this point of view, we also believe that our rounding scheme is of independent interest when other problems with hard constraints are considered.

Our technical tool for establishing the approximation guarantee for the final rounding step is a separation lemma that certifies the conditional existence of a strong partition on the vertices, given by solutions that are basic feasible for the LPs our algorithm is working with. The presented separation lemma can be seen as a recast of the well-established Carathéodory's theorem [29] in the language of VC-HC polytope. Intuitively, it says that: A vertex-to-edge correspondence with strong separation guarantee exists at places where a certain type of linear constraint is inactive. Similar ideas have been used, e.g., [6], in the usage of the multi-set multi-cover polytope [28], which by formulation is a vertex-to-vertex matching structure, to extract a vertex-to-vertex matching. When the VC-HC polytope is considered, however, we have an essentially different entangling relation to deal with, in particular, relation between the vertices $V$, the edges $E$, and their cross-products $E \times V$. Our separation lemma provides an interpretation from the perspective of Carathéodory's theorem to this situation: By properly removing the entangling $E \times V$ relation, a strong and meaningful structural guarantee can be extracted for our rounding approach.

The rest of this paper is organized as follows. In Section 2 we formally define VC-HC and the extended LP relaxation we will be using throughout this paper. In Section 3 we introduce the key notion of edge-folding that motivates our iterative rounding approach. We present our tight approximation algorithm in Section 4 and establish the approximation guarantee in Section 5.

### 1.2 Further Related Work

In the following we briefly summarize other related results for vertex cover (VC), CVC, and VC-HC with relaxed constraints.

**Soft capacitated covering.** For vertex cover, it is known that a $f$-approximation can be obtained by both LP rounding and LP duality [4, 17] for hypergraphs. Khot and Regev [23] showed that, by assuming the unique game conjecture (UGC), approximating this problem to a ratio better than $f - \epsilon$ is NP-hard for any $\epsilon > 0$ and $f \geq 2$.

The soft capacitated vertex cover problem was first introduced in the notable work of Guha et al [15], in which a 2-approximation was presented using primal-dual approach. This result extends to $f$-approximation for hypergraphs. Gandhi et al. [10] provided another 2-approximation via dependent rounding. Kao et al [19–21] considered capacitated dominating set, an alternative notion of capacitated covering, and presented a series of results studying the complexity and approximability of this problem when different classes of graphs are considered. Special cases and variations of this problem were also considered independently [8, 9, 27].

**VC-HC with relaxed constraints.** For partial VC-HC, which aims at covering a given number of edges, Cheung et al [6] provided a $(1+\epsilon)(2f+2)$-approximation, based on a reduction to $(f+1)$-hypergraphs and their $2f$-approximation for VC-HC presented in the same work.

Gandhi et al [12] considered weighted VC-HC with relaxed multiplicity constraints and presented an augmented $(2,2)$-covering[1] for normal graphs. Grandoni et al. [14] considered unit vertex multiplicity, i.e., $m(v)=1$, and presented an augmented $(f, f^2)$-covering for hypergraphs. These results were improved by Kao et al. [22] to augmented $\big(s, (1+\frac{1}{s-1})(f-1)\big)$-covering for any $s \geq 2$, using an extended dual-fitting approach.

### 1.3 Other hard-capacitated problems.

In recent years, significant progress have been made on hard-capacitated problems, including VC-HC, Capacitated facility location (CFL), Capacitated $k$-center, and $k$-median. We summarize the current progress and recent breakthroughs for the respective problems in Table 1, using a condensed format. We also refer the reader to the references therein for further details.

<table>
<tr><td rowspan="3">VC<br><br>( Set Cover )</td><td rowspan="3">open<br>cost</td><td>CVC</td><td colspan="2">$f$-approx via primal-dual [15]</td></tr>
<tr><td rowspan="2">VC-HC</td><td>weighted</td><td>unweighted</td></tr>
<tr><td>set-cover-hard [7]</td><td>$f$-approx via<br>iterative partial rounding</td></tr>
<tr><td colspan="3">logarithmic approx. [7, 30]</td></tr>
<tr><td rowspan="2">Facility<br>Location</td><td rowspan="2">open<br>+<br>assign</td><td>UFL</td><td colspan="2">1.488-approx [24],<br>NP-hard to approx within $(1.463-\epsilon)$ [16]</td></tr>
<tr><td>CFL</td><td colspan="2">5-approx via local search [3],<br>$O(1)$-approx, LP-rounding via MFN-relaxation [2]</td></tr>
<tr><td rowspan="2">$k$-center</td><td rowspan="3">assign<br>cost</td><td rowspan="2">max min</td><td colspan="2">2-approximation [18]</td></tr>
<tr><td>Hard-capd.</td><td>9-approximation [1]</td></tr>
<tr><td rowspan="2">$k$-median</td><td rowspan="2">avg</td><td colspan="2">uncapacitated,<br>$(2.675+\epsilon)$-approx via pseudo-approx [5, 26]</td></tr>
<tr><td colspan="2">uniform hard-capacitated,<br>$e^{O(1/\epsilon^2)}$-bi-approx, using $(1+\epsilon)\cdot k$ facilities [25]</td></tr>
</table>

Table 1: A summary on related hard-constrained problems.

[1] An augmented $(\alpha, \beta)$-covering is a solution that violates the multiplicity limit by a factor $\alpha$ and has a cost factor $\beta$ to the optimal LP solution for VC-HC.

## 2 Problem Statement and LP Relaxation

In this section we define VC-HC formally and introduce the extended LP relaxation we will be using in this paper.

We begin with basic graph notations. Throughout this paper, we use $G = (V, E)$ to denote a hypergraph with vertex set $V$ and edge set $E \subseteq 2^V$. Note that, each hyperedge $e \in E$ is represented by the set of its incident vertices, which is a vertex subset of $V$. We use $f := \max_{e \in E} |e|$ to denote the size of the largest hyperedge in the considered graph $G$.

For any edge subset $\mathcal{E} \subseteq E$, we use $\mathcal{E}[v]$ to denote the set of edges in $\mathcal{E}$ that are incident to the vertex $v$. Formally, $\mathcal{E}[v] := \big\{\, e : \ e \in \mathcal{E} \text{ such that } v \in e \,\big\}$. This definition extends to sets of vertices, i.e., for any $U \subseteq V$, $\mathcal{E}[U] := \bigcup_{v \in U} \mathcal{E}[v] = \big\{\, e : \ e \in \mathcal{E} \text{ such that } e \cap U \neq \emptyset \,\big\}$.

**Minimum Vertex Cover with Hard Capacities (VC-HC).** In this problem we are given a hypergraph $G = (V, E)$, where each $e \in E$ is associated with a demand $d_e$ and each $v \in V$ is associated with a capacity $c_v$ and an (integral) available multiplicities $m_v$.

A solution to this problem consists of an assignment $h \colon E \times V \to \mathbb{R}^+ \cup \{0\}$, where $h_{e,v}$ denotes the fraction of the edge $e$ that is assigned to the vertex $v$. The multiplicity given by $h$ is denoted $x_v^{(h)} := \left\lceil \sum_{e \in E[v]} (d_e \cdot h_{e,v})/c_v \right\rceil$. The assignment $h$ is said to be *feasible* if $\sum_{v \in e} h_{e,v} = 1$ for all $e \in E$ and $x_v^{(h)} \leq m_v$ for all $v \in V$.

Given an instance $\Pi = (V, E, \mathbf{c}, \mathbf{m}, \mathbf{d})$ as described above, the problem VC-HC is to find a feasible assignment $h$ such that $\sum_{v \in V} x_v^{(h)}$ is minimized. Without loss of generality, we assume that the input graph $G$ admits a feasible assignment since this condition can be checked via a max-flow computation.

We also remark that, when $d_e$ and $c_v$ are integer-valued for all $e \in E$, $v \in V$, by the integrality of b-matching polytope, any fractional assignment can be turned into an integral assignment, i.e., $d_e \cdot h_{e,v} \in \mathbb{Z}^{\geq 0}$ for all $e, v$, using a standard integer flow computation.

**Extended LP relaxation for VC-HC.** Given an instance $\Pi = (V, E, \mathbf{c}, \mathbf{m}, \mathbf{d})$ of VC-HC and an additional lower-bound $\boldsymbol{\ell}$, where $\mathbf{0} \leq \boldsymbol{\ell} \leq \mathbf{m}$, on the multiplicity of the vertices, we consider the following LP relaxation, with $\Psi = (V, E, \boldsymbol{\ell}, \mathbf{c}, \mathbf{m}, \mathbf{d})$ being a parameter tuple:

LP($\Psi$)

$$\min_{(\mathbf{x},\mathbf{h})} \ \sum_{v \in V} x_v \qquad \text{s.t.} \ \ (\mathbf{x}, \mathbf{h}) \in \mathbf{Q}(\Psi)$$

The Polytope $\mathbf{Q}(\Psi)$:

$$\sum_{v \in e} h_{e,v} = 1, \qquad \forall e \in E \tag{1a}$$

$$\sum_{e \in E[v]} d_e \cdot h_{e,v} \leq c_v \cdot x_v, \qquad \forall v \in V \tag{1b}$$

$$\ell_v \leq x_v \leq m_v, \qquad \forall v \in V \tag{1c}$$

$$0 \leq h_{e,v} \leq x_v, \qquad \forall e \in E, \ v \in e \tag{1d}$$

Since each of the variables $h_{e,v}$ and $x_v$ is bounded from both below and above, we know that $\mathbf{Q}(\Psi)$ is indeed a polytope, and the reference to its extreme points[2] is well-defined.

[2]Recall that $p \in \mathbf{Q}(\Psi)$ is an extreme point if it is not in the interior of any line segment contained in $\mathbf{Q}(\Psi)$, i.e., $p = \lambda r + (1 - \lambda)s$ for some $0 < \lambda < 1$ implies that either $r \notin \mathbf{Q}(\Psi)$ or $s \notin \mathbf{Q}(\Psi)$.

Throughout this paper, for any given instance $\Pi$ of VC-HC, a number of different parameter tuples will be considered. However, since $V$, $\mathbf{m}$, and $\mathbf{d}$ will remain the same in every considered tuple, we will simply use $(E, \boldsymbol{\ell}, \mathbf{c})$ to denote the parameter tuple $\Psi$ for the considered LP.

# 3 The VC Constraint and Edge Folding

Our approximation algorithm is inspired by a simple one-edge example and a review on the natural LP relaxation, in particular, the constraint (1d) which connects VC-HC to the classical vertex cover (VC) and reveals the uncapacitated problem core of VC-HC. In the following we elaborate this idea in further detail.

Consider LP($\Psi$) with $\Psi = (E, \mathbf{0}, \boldsymbol{c})$, i.e., the original natural LP relaxation. It is well-understood that, although the constraints (1a) to (1c) together with $h_{e,v} \geq 0$ for all $e, v$ already formulate VC-HC, it does not yield a solution with bounded integrality gap since it allows vertices to take arbitrarily small multiplicities. This is illustrated by the one-edge example shown in Figure 1. The constraints (1a) to (1c) can be satisfied by setting $h_{\tilde{e},\tilde{v}} = 1$ and $x_{\tilde{v}} = \epsilon \approx 0$. As a result, the fractional solutions can deviate arbitrarily from the integral solution for which we would have $x_{\tilde{v}} = 1$.

This is how constraint (1d) is introduced to bound integrality gap. With this constraint in place, the value of $x_{\tilde{v}}$ is lower-bounded by $h_{\tilde{e},\tilde{v}}$. This ensures that the fractional solution in this example will agree with its integral solution. One explanation to this phenomenon is that constraint (1d) connects this relaxation to the classical VC problem — a problem for which we have extensively stuidied and for which we understand very well.

Figure 1:

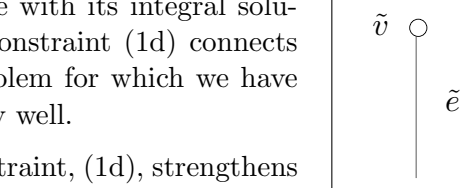


**Our key observation** here is that, while the VC constraint, (1d), strengthens the LP relaxation, it also hints on how the rounding can be done locally when it is tight. Suppose that $h_{e',v} = x_v$ holds for some $e'$ and $v$. Then from (1b) we know that

$$d_{e'} \cdot h_{e',v} \leq \sum_{e:\, v \in e} d_e \cdot h_{e,v} \leq c_v \cdot x_v,$$

and it follows that $d_{e'} \leq c_v$ since $h_{e',v} = x_v$.

This suggests that: (1) The existence of $e'$ is vital in support for the value of $x_v$, despite the fact that $d_{e'}$ can be substantially smaller than $c_v$. (2) If the vertex $v$ is to be rounded up eventually, then we know that its capacity will be sufficient for covering $d_{e'}$. This motivates the key notion of *structural support* and the *edge-folding* operation.

**Definition 1** (Supporting edge)**.** Let $\Psi = (\hat{E}, \boldsymbol{\ell}, \boldsymbol{c})$ be a parameter tuple and $\hat{p} = (\hat{x}, \hat{h})$ be a feasible solution for LP($\Psi$). For any $e \in \hat{E}$ and $v \in e$, we say that edge $e$ (structurally) *supports* vertex $v$ in $\hat{p}$ if $0 < \hat{h}_{e,v} = \hat{x}_v$.

**Definition 2** (Edge folding)**.** Let $\Psi = (\hat{E}, \boldsymbol{\ell}, \boldsymbol{c})$ be the current tuple, $\hat{p} = (\hat{x}, \hat{h})$ be a feasible solution for LP($\Psi$), and $e$ be an edge that supports a vertex $v$ in $\hat{p}$. By *folding $e$ into $v$*, we:

(i) Remove $e$ from $\hat{E}$ and decrease $c_v$ by $d_e$.

(ii) Impose a new constraint $x_v \geq \hat{x}_v$ to LP($\Psi$) by setting $\ell_v = \hat{x}_v$.

Note that, from the observation described above, folding supporting edges into the supported vertices results in no loss in the feasibility and approximation guarantee of the final solution if we know in advanced that the supported vertices are to be rounded-up eventually.

# 4 Iterative Partial Rounding for VC-HC

The observation in the previous section motivates our iterative partial rounding approach, which focuses on tackling the VC constraint – an uncapacitated problem core of VC-HC, using the edge-folding operation. At a high-level description, it is done as follows. We start with the initial parameter tuple $\Psi = (E, \mathbf{0}, \boldsymbol{c})$.

1. Solve $\mathrm{LP}(\Psi)$ for a *basic optimal solution* $p = (\boldsymbol{x}, \boldsymbol{h})$.
   Let $I := \{v \colon 0 < x_v < \frac{1}{f}\}$.
2. If there exists an edge $e$ that structurally supports some $v \notin I$ in $p$,

   *then*
   - Fold $e$ into $v$.
   - Restart Step 1.

   *otherwise*
   - Round up $\boldsymbol{x}$ and stop.

In contrast to previous approaches, e.g., [6, 7, 28], which round up every vertex with decent value at once, we only deal with those that are structurally supported. Furthermore, instead of rounding up the variables aggressively to one, we round them partially to a fractional value by imposing stronger lower bound constraints on them. This nature allows previously non-supporting edges to become supporting as the algorithm iterates. The key feature of this approach is that it offers a series of LPs with non-increasing objective values which preserves the optimality of the initial LP for the final rounding step.

Let $\hat{\Psi} = (\hat{\mathcal{E}}, \hat{\boldsymbol{\ell}}, \hat{\boldsymbol{c}})$ denote the parameter tuple when the algorithm enters the final rounding step. Let $\boldsymbol{h}'$ denote the $\{0,1\}$-assignment function for $E \setminus \hat{\mathcal{E}}$ that indicates the vertex which each edge $e \in E \setminus \hat{\mathcal{E}}$ is folded into. In particular, for each $e \in E \setminus \hat{\mathcal{E}}$ and $v \in e$, $h'_{e,v}$ is set to 1 if $e$ is folded into $v$ and 0 otherwise. Let $\hat{p} = (\hat{x}, \hat{h})$ denote the basic optimal solution computed for $\hat{\Psi}$.

The final output $(\boldsymbol{x}^*, \boldsymbol{h}^*)$ is defined as follows. For any $v \in V$ and $e \in E[v]$, let

$$x^*_v := \left\lceil \hat{x}_v \right\rceil \quad \text{and} \quad h^*_{e,v} := \begin{cases} \hat{h}_{e,v}, & \text{if } e \in \hat{\mathcal{E}}, \\ h'_{e,v}, & \text{otherwise.} \end{cases}$$

Then $(\boldsymbol{x}^*, \boldsymbol{h}^*)$ is output as the solution.

A pseudo-code of this algorithm is provided in Figure 4 in Page 10 for further reference. We remark that, since this approach is essentially insensitive to multiple foldings, in the actual algorithm we will fold every supporting edge, provided that it supports some $v \notin I$. Furthermore, ties are broken arbitrarily if an edge supports multiple vertices outside $I$.

Let Tight-VC-HC denote the above algorithm. Our tight approximation result is stated in the following theorem:

**Theorem 2.** On any instance $\Pi = (V, E, \boldsymbol{c}, \boldsymbol{m}, \boldsymbol{d})$ of VC-HC with maximum edge size $f \geq 2$, algorithm Tight-VC-HC outputs an $f$-approximation $(\boldsymbol{x}^*, \boldsymbol{h}^*)$ in polynomial time.

From the usage of edge-folding operation in the algorithm and the intuition provided in Section 3 for the VC constraints, it is not difficult to see that:

- Algorithm Tight-VC-HC terminates in $O(|E|)$ rounds.
- The output $(\boldsymbol{x}^*, \boldsymbol{h}^*)$ is indeed feasible for the initial $\mathrm{LP}\big((E, \mathbf{0}, \boldsymbol{c})\big)$, and $\boldsymbol{x}^*$ is integral.

Due to the page limit, the proof for the feasibility of algorithm Tight-VC-HC is provided in Section B.1 for further reference. In order to fully-verify the detail provided in Section B.1, we also refer the reader to the notations and properties given in the beginning of Section B. In the remaining of this paper, we will describe how the approximation guarantee is established.

# 5 Approximation Guarantee

In this section we establish our approximation guarantee. The argument we use builds on the fact that $p = (\boldsymbol{x}, \boldsymbol{h})$ is a basic feasible solution for the extended LP relaxation LP($\Psi$). In Section 5.1 we first define the notions that will be used to capture the structural properties given by our partial rounding approach. In Section 5.2 we formally state our separation lemma, which is our main technical tool for establishing the approximation guarantee, and discuss the intuitions behind. In Section 5.3, we will describe how the technical separation lemma is used to establish the tight approximation guarantee.

## 5.1 Notion and Definitions for Extremality

In this section we introduce notions that are required to state our separation lemma for the polytope $\mathbf{Q}(\Psi)$. Let $p = (\mathbf{x}, \mathbf{h})$ be a point in $\mathbf{Q}(\Psi)$, where $\Psi = (\hat{E}, \boldsymbol{\ell}, \boldsymbol{c})$ is a parameter tuple. The following terminologies are defined with reference to point $p$ and parameter tuple $\Psi$.

**Definition 3** (Non-extremal vertex)**.** For any $v \in V$, we say that $v$ is *non-extremal* if $\ell_v < x_v < m_v$. Otherwise $v$ is said to be *extremal*.

Intuitively, a vertex is non-extremal if and only if $x_v$ is not constrained by (1c) in $\mathbf{Q}(\Psi)$. For any $e \in \hat{E}$ and $v \in e$, we say that $v$ is an *active end* of $e$ if $h_{e,v} > 0$.

**Definition 4** (Active subedge and active edge sets)**.** For any $e \in \hat{E}$, we define the *active subedge* of $e$, denoted $e_h^{\text{actv}}$, to be the set of its active ends, i.e.,

$$e_h^{\text{actv}} := \Big\{ v : \ v \in e, \ \ h_{e,v} > 0 \Big\}.$$

For any edge subset $\mathcal{E} \subseteq \hat{E}$, we extend the above definition and use $\mathcal{E}_h^{\text{actv}} := \Big\{ e_h^{\text{actv}} : \ e \in \mathcal{E} \Big\}$ to denote the set of active subedges of the edges in $\mathcal{E}$.

Intuitively, the active subedge corresponds to the set of vertices to which the demand of an edge is assigned in $\boldsymbol{h}$. Note that, since $p \in \mathbf{Q}(\Psi)$, it follows from constraint (1a) that $e_h^{\text{actv}} \neq \emptyset$ for all $e \in \hat{E}$. Recall that we also use $\mathcal{E}[U]$ for a vertex subset $U$ of $V$ to denote the set of edges in $\mathcal{E}$ that are incident to some vertex in $U$. Given the definition of active edge sets, $\mathcal{E}_h^{\text{actv}}[U]$ is then used to denote the set of incident (sub)edges $U$ has in $\mathcal{E}_h^{\text{actv}}$, i.e.,

$$\mathcal{E}_h^{\text{actv}}[U] := \Big\{ \tilde{e} : \ \tilde{e} \in \mathcal{E}_h^{\text{actv}} \text{ such that } \tilde{e} \cap U \neq \emptyset \Big\}.$$

Intuitively, this corresponds to the set of active subedges from $\mathcal{E}$ that intersect $U$. Note that it is a collection of subedges rather than the original edges.

**Definition 5** (Supporting sets and supported sets)**.** For a vertex subset $U \subseteq V$, we say that

- $U$ is *supporting* if there exists an edge $e \in \hat{E}$ with $e_h^{\text{actv}} \cap U \neq \emptyset$ and a vertex $v \in e_h^{\text{actv}} \setminus U$ such that $e$ supports $v$.
- $U$ is *supported* if there exists $e \in \hat{E}$ and $v \in e_h^{\text{actv}} \cap U$ such that $v$ is supported by $e$.

The definition of supporting sets and supported sets may seem unnatural in the beginning. However, consider the set $I := \big\{v \colon 0 < x_v < \frac{1}{f}\big\}$ of vertices with small fractional values. In our partial rounding approach, we iteratively remove the support from vertices with decent fractional values, i.e., those not in $I$, and when the algorithm terminates, we know that the set of small vertices is not supporting while the set of vertices with large fractional values, e.g., those with $x_v > 1$, is not supported.

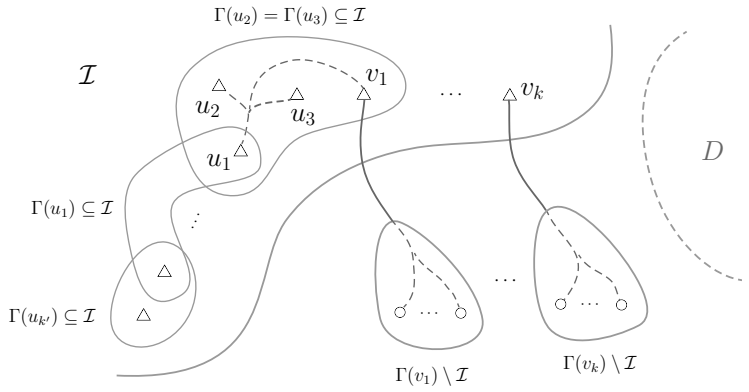


Figure 2: A depiction on the set $\mathcal{I}$ and the subedges given by $\Gamma$. $\{u_i\}_{1 \le i \le k'} \subseteq \mathcal{I}$ denote vertices with $\Gamma(u_i) \subseteq \mathcal{I}$ and $\{v_i\}_{1 \le i \le k}$ denote vertices with $\Gamma(v_i) \not\subseteq \mathcal{I}$, i.e., $\Gamma(v_i) \setminus \mathcal{I} \neq \emptyset$.

### 5.2 A Separation Lemma for Polytope $\mathbf{Q}(\Psi)$

We formally state our structural separation lemma, which states that a strong partition exists when the edge-to-vertex supporting relations are properly eliminated:

**Lemma 3** (Existence of a strong partition)**.** Let $p = (\mathbf{x}, \mathbf{h})$ be an extreme point of polytope $\mathbf{Q}(\Psi)$. For any disjoint sets $\mathcal{I}$, $\mathcal{D}$ of non-extremal vertices, if $\mathcal{I}$ is not supporting and $\mathcal{D}$ is not supported, then there exists a mapping [3]

$$\Gamma : \ \mathcal{I} \ \mapsto \ E_h^{\text{actv}}[\mathcal{I}] \setminus E_h^{\text{actv}}[\mathcal{D}]$$

such that the following hold for any $v \in \mathcal{I}$:

1. (reflexive) $v \in \Gamma(v)$.
2. (closed under intersection) $\Gamma(u) \cap \Gamma(v) \subseteq \mathcal{I}$ for any $u \in \mathcal{I} \setminus \{v\}$.

Given sets $\mathcal{I}$ and $\mathcal{D}$ as stated in the prerequisite, Lemma 3 says that, there exists a mapping $\Gamma$ such that, for each $v \in \mathcal{I}$, the active subedge $\Gamma(v)$ contains the vertex $v$ but no vertices in $\mathcal{D}$. Furthermore, the intersection of these active subedges can happen only inside $\mathcal{I}$, i.e., they are mutually disjoint outside the set $\mathcal{I}$. See also Figure 2 for an illustration.

Note that, the statement of Lemma 3 does not exclude the possibility that $\Gamma(v) \subseteq \mathcal{I}$ and the mapping $\Gamma$ may not necessarily be injective. In fact, it depends on the choice of $\mathcal{I}$, since we make no prior assumption on the edge-vertex supporting relation inside $\mathcal{I}$. When all non-extremal vertices are selected into $\mathcal{I}$, then Figure 2 is what we can expect.

With an appropriate choice of $\mathcal{I}$, e.g., one that ensures $\tilde{e} \not\subseteq \mathcal{I}$ for all $\tilde{e} \in E_h^{\text{actv}}[\mathcal{I}]$, however, the injective property of $\Gamma$ will follow from the two properties stated in Lemma 3. Notice that the set of vertices with small fractional values is exactly one of such choices. In Section 5.3 we will illustrate how the mapping $\Gamma$ is used to obtain a tight approximation result.

In the following we discuss the technique we use to build this lemma and the intuitions behind. A formal proof to this lemma is provided in Section C for further reference.

**Extreme points of the VC-HC polytope.** Lemma 3 can be seen as a recast of the well-established Carathéodory's theorem [29] in the language of VC-HC polytope. Intuitively, it says that: A vertex-to-edge correspondence with strong separation guarantee exists at places where the the VC constraint is inactive.

[3] We would like to refer the reader to the previous section for the definitions of non-extremal vertices, active subedges, and supporting(-ed) sets in order to fully-access the structural message in this lemma.

Below we further elaborate this idea. Let $p = (\boldsymbol{x}, \boldsymbol{h})$ be an extreme point of the VC-HC polytope and $\mathrm{Var}(\mathcal{I}, \mathcal{D})$ denote the set of variables that are related to $\mathcal{I}$, $\mathcal{D}$ and the edges in $E[\mathcal{I}]$. By Carathéodory's theorem, the variables in $\mathrm{Var}(\mathcal{I}, \mathcal{D})$ are constrained by a set, with the same cardinality, of linearly independent inequalities that hold with equality.

This gives a one-to-one (injective) correspondence between the variables and the constraints. From the assumption that $\mathcal{I}$ is non-extremal and not supporting, it follows that:

(i) The variable $x_v$ *must* be constrained by either (1b) or (1d), for all $v \in \mathcal{I}$.

(ii) The variable $h_{e,u}$, where $h_{e,u} \in \mathrm{Var}(\mathcal{I}, \mathcal{D})$ such that $e \cap \mathcal{I} \neq \emptyset$ and $u \notin \mathcal{I}$, *can only* be constrained by (1a) and (1b).

From (i) and the fact that the constraints extracted by Carathéodory's theorem are linearly independent, we show that for each $v \in \mathcal{I}$, an unique edge constraint, i.e., (1a), can be identified. This gives the edge for $\Gamma(v)$ to be defined. Then the main claim of this lemma is guaranteed by (ii) and the injective property of the variable-to-constraint mapping.

To see that the vertices in $\mathcal{D}$ can be excluded from the subedges defined above, it suffices to observe that for any $v \in \mathcal{D}$, the only constraint for variable $x_v$ to be constrained is (1b), i.e., $v$ itself, since it is non-extremal and not supported by definition. Therefore it cannot be included in any of these active subedges since their corresponding constraints have already been occupied. The complete proof of this lemma is provided in Section C for further reference.

### 5.3 Establishing the Approximation Guarantee

In this section we describe how Lemma 3 is used to obtain our tight approximation guarantee. Let $\hat{p} = (\hat{x}, \hat{h})$ denote the basic optimal solution computed in the final iteration of the algorithm. Consider the two sets

$$I := \left\{ v : \ v \in V, \ \ 0 < \hat{x}_v < \frac{1}{f} \right\} \quad \text{and} \quad D := \left\{ v : \ v \in V, \ \ 1 < \hat{x}_v < m_v \right\}.$$

It follows that $e_h^{\mathrm{actv}} \not\subseteq I$ for all $e$ such that $e_h^{\mathrm{actv}} \cap I \neq \emptyset$. Therefore condition (ii) of Lemma 3 implies that $\Gamma$ will be injective, i.e., $\Gamma(u) \neq \Gamma(v)$ whenever $u \neq v$, and thereby defines an equivalence relation in $V \setminus I$ with respect to $I$, witnessed by the active subedge $\Gamma(v)$ for each $v \in I$. See also Figure 3 for an illustration of this partition.

Consider the active subedge $\Gamma(v)$ for some $v \in I$. By Lemma 3, we know that $\Gamma(v)$ contains no vertices in $D$, since $\Gamma(v) \notin E_h^{\mathrm{actv}}[D]$. It follows that, for each $u \in \Gamma(v)$, either $\hat{x}_u = m_u$ or $\hat{x}_u \leq 1$ will hold. For the former case, the rounding cost of $v$ can be absorbed by $u$, since we have $m_u \geq 1$ and $f \geq 2$. Below we consider the case that $\hat{x}_u \leq 1$ for all $u \in \Gamma(v)$.

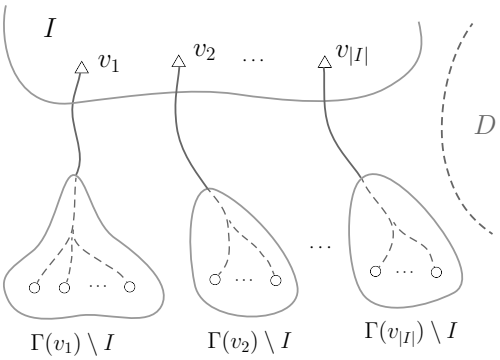


Figure 3: A partition of vertices in $V \setminus I$, witnessed by the active subedges $\{\Gamma(v_i)\}_{1 \leq i \leq |I|}$

Suppose that $\hat{x}_u \le 1$ for all $u \in \Gamma(v)$. Let $e$ be an edge whose active subedge is exactly $\Gamma(v)$. From the constraint (1a) we have $\sum_{u\in\Gamma(v)} \hat{h}_{e,u} = 1$. Since we know that $\hat{x}_u$ is small for all $u \in \Gamma(v) \cap I$, it follows that

$$\hat{x}_v + \sum_{u\in\Gamma(v)\setminus I} \hat{x}_u \;\ge\; 1 - \sum_{u\in(\Gamma(v)\cap I)\setminus\{v\}} \hat{x}_u \;\ge\; 1 - \frac{1}{f}\cdot\Big|\,(\Gamma(v)\cap I)\setminus\{v\}\,\Big| \;\ge\; \frac{1}{f}\cdot\Big|\,\big(\Gamma(v)\setminus I\big)\cup\{v\}\,\Big|,$$

where in the last inequality we use the fact that $\big|(\Gamma(v)\cap I)\setminus\{v\}\big| = \big|\Gamma(v)\big| - \big|(\Gamma(v)\setminus I)\cup\{v\}\big|$ and $\big|\Gamma(v)\big| \le |e| \le f$. Therefore the rounding cost incurred by $v$ and vertices in $\Gamma(v)\setminus I$ can be bounded within $f\cdot\big(\hat{x}_v + \sum_{u\in\Gamma(v)\setminus I}\hat{x}_u\big)$, for any $f \ge 2$.

We remark that the exclusion of $D$ from the image of $\Gamma$ is crucial in obtaining a tight approximation for $f = 2$, i.e., the normal graphs. The reason is that, rounding up medium-sized vertices, e.g., $u$ with $1 < \hat{x}_u < 1 + 1/2$, already results in a factor of 2, rendering them unable to absorb additional rounding cost incurred by small vertices. A formal proof for the approximation guarantee is provided in Section B.2 for further reference.


## Acknowledgement

The author would like to thank Kai-Min Chung, Herbert Yu, and the anonymous referees for their very helpful comments on the presentation of this work.


**Input:** an instance $\Pi = (V, E, \mathbf{c}, \mathbf{m}, \mathbf{d})$ for VC-HC.
**Output:** an $f$-approximation $(\boldsymbol{x}^*, \boldsymbol{h}^*)$ for $\Pi$, where $f = \max_{e\in E}|e|$.

Step 1. Let $\mathcal{E} = E$.
Set $h'_{e,v} = 0$ for all $e \in E$, $v \in e$ and $\begin{cases}\ell_v = 0,\\ c'_v = c_v,\end{cases}$ for all $v \in V$.

Step 2. Repeat the following procedure:

(2.a) Solve $\mathrm{LP}\Big((\mathcal{E}, \boldsymbol{\ell}, \boldsymbol{c'})\Big)$ for a basic optimal solution $(\boldsymbol{x}, \boldsymbol{h})$.

(2.b) Let $I := \Big\{\, v :\ v \in V,\ \ 0 < x_v < \frac{1}{f} \,\Big\}$.

(2.c) For each $e \in \mathcal{E}$, define $T(e) := \Big\{\, v :\ v \in e\setminus I,\ \ 0 < h_{e,v} = x_v \,\Big\}$.

(2.d) Go to Step 3 if $\bigcup_{e\in\mathcal{E}} T(e) = \emptyset$.

(2.e) For each $e \in \mathcal{E}$ such that $T(e) \ne \emptyset$, do

i. Pick an *arbitrary* $v \in T(e)$ and set $h'_{e,v} = 1$.

ii. For *all* $v \in T(e)$, decrease $c'_v$ by $d_e$ and set $\ell_v = x_v$.

iii. Remove $e$ from $\mathcal{E}$.

Step 3. For any $v \in V$ and any $e \in E[v]$, define

$$x^*_v := \Big\lceil x_v \Big\rceil \quad\text{and}\quad h^*_{e,v} := \begin{cases} h_{e,v}, & \text{if } e \in \mathcal{E},\\ h'_{e,v}, & \text{otherwise.}\end{cases}$$

Output $(\boldsymbol{x}^*, \boldsymbol{h}^*)$.

Figure 4: Algorithm Tight-VC-HC.

## References


[1] Hyung-Chan An, Aditya Bhaskara, Chandra Chekuri, Shalmoli Gupta, Vivek Madan, and Ola Svensson. Centrality of trees for capacitated $k$-center. *Math. Program.*, 154(1-2):29–53, 2015.

[2] Hyung-Chan An, Mohit Singh, and Ola Svensson. LP-based algorithms for capacitated facility location. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 256–265, 2014.

[3] Manisha Bansal, Naveen Garg, and Neelima Gupta. A 5-approximation for capacitated facility location. In *Proceedings of the 20th Annual European Conference on Algorithms*, ESA'12, pages 133–144, Berlin, Heidelberg, 2012. Springer-Verlag.

[4] Reuven Bar-Yehuda and Shimon Even. A linear-time approximation algorithm for the weighted vertex cover problem. *Journal of Algorithms*, 2(2):198 – 203, 1981.

[5] Jaroslaw Byrka, Thomas Pensyl, Bartosz Rybicki, Aravind Srinivasan, and Khoa Trinh. An improved approximation for $k$-median, and positive correlation in budgeted optimization. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, pages 737–756. SIAM, 2015.

[6] W.-C. Cheung, M. Goemans, and S. Wong. Improved algorithms for vertex cover with hard capacities on multigraphs and hypergraphs. SODA'14, 2014.

[7] Julia Chuzhoy and Joseph Naor. Covering problems with hard capacities. *SIAM Journal on Computing*, 36(2):498–515, August 2006.

[8] Marek Cygan, Marcin Pilipczuk, and Jakub Onufry Wojtaszczyk. Capacitated domination faster than $O(2^n)$. SWAT'10, pages 74–80, 2010.

[9] M. Dom, D. Lokshtanov, S. Saurabh, and Y. Villanger. Capacitated domination and covering: A parameterized perspective. IWPEC'08, pages 78–90, 2008.

[10] R. Gandhi, S. Khuller, S. Parthasarathy, and A. Srinivasan. Dependent rounding in bipartite graphs. In *Foundations of Computer Science, 2002. Proceedings. The 43rd Annual IEEE Symposium on*, pages 323–332, 2002.

[11] Rajiv Gandhi, Eran Halperin, Samir Khuller, Guy Kortsarz, and Aravind Srinivasan. An improved approximation algorithm for vertex cover with hard capacities. *J. Comput. Syst. Sci.*, 72:16–33, February 2006.

[12] Rajiv Gandhi, Samir Khuller, Srinivasan Parthasarathy, and Aravind Srinivasan. Dependent rounding and its applications to approximation algorithms. *J. ACM*, 53(3):324–360, May 2006.

[13] Rajiv Gandhi, Samir Khuller, and Aravind Srinivasan. Approximation algorithms for partial covering problems. *Journal of Algorithms*, 53(1):55–84, October 2004.

[14] F. Grandoni, J. Könemann, A. Panconesi, and M. Sozio. A primal-dual bicriteria distributed algorithm for capacitated vertex cover. *SIAM J. Comput.*, 38(3), 2008.

[15] Sudipto Guha, Refael Hassin, Samir Khuller, and Einat Or. Capacitated vertex covering. *Journal of Algorithms*, 48(1):257–270, August 2003.

[16] Sudipto Guha and Samir Khuller. Greedy strikes back: Improved facility location algorithms. In *Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '98, pages 649–657, Philadelphia, PA, USA, 1998. Society for Industrial and Applied Mathematics.

[17] Dorit S. Hochbaum. Approximation algorithms for the set covering and vertex cover problems. *SIAM Journal on Computing*, 11(3):555–556, 1982.

[18] Dorit S. Hochbaum and David B. Shmoys. A best possible heuristic for the $k$-center problem. *Math. Oper. Res.*, 10(2):180–184, May 1985.

[19] Mong-Jen Kao. *An Algorithmic Approach to Local and Global Resource Allocations*. PhD thesis, National Taiwan University, 2012.

[20] Mong-Jen Kao, Han-Lin Chen, and D.T. Lee. Capacitated domination: Problem complexity and approximation algorithms. *Algorithmica*, November 2013.

[21] Mong-Jen Kao, Chung-Shou Liao, and D. T. Lee. Capacitated domination problem. *Algorithmica*, 60(2):274–300, June 2011.

[22] Mong-Jen Kao, Hai-Lun Tu, and D.T. Lee. An $O(f)$ bi-approximation for weighted capacitated covering with hard capacities. manuscript, 2015.

[23] Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within $2-\epsilon$. *Journal of Computer and System Sciences*, 74(3):335–349, May 2008.

[24] Shi Li. A 1.488 approximation algorithm for the uncapacitated facility location problem. In *Proceedings of the 38th International Conference on Automata, Languages and Programming - Volume Part II*, ICALP'11, pages 77–88, Berlin, Heidelberg, 2011. Springer-Verlag.

[25] Shi Li. On uniform capacitated $k$-median beyond the natural lp relaxation. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, pages 696–707. SIAM, 2015.

[26] Shi Li and Ola Svensson. Approximating $k$-median via pseudo-approximation. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 901–910, New York, NY, USA, 2013. ACM.

[27] Mathieu Liedloff, Ioan Todinca, and Yngve Villanger. Solving capacitated dominating set by using covering by subsets and maximum matching. WG'10, pages 88–99, 2010.

[28] Barna Saha and Samir Khuller. Set cover revisited: Hypergraph cover with hard capacities. ICALP'12, pages 762–773, 2012.

[29] Alexander Schrijver. *Theory of Linear and Integer Programming*. John Wiley & Sons, Inc., New York, NY, USA, 1986.

[30] L.A. Wolsey. An analysis of the greedy algorithm for the submodular set covering problem. *Combinatorica*, 2(4):385–393, 1982.

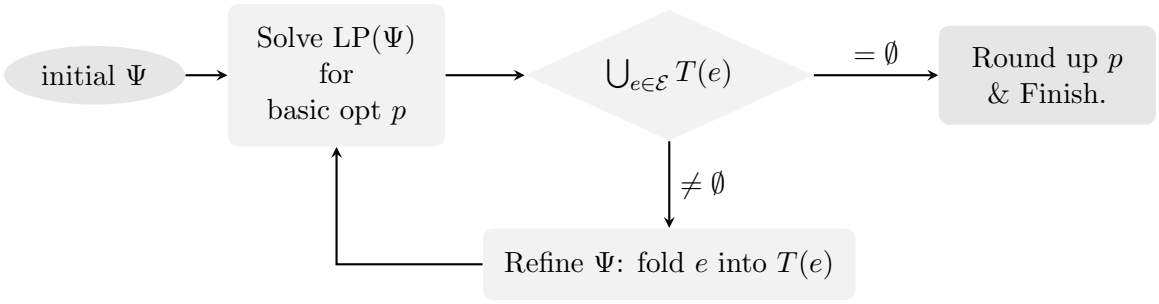


Figure 5: Overview of Algorithm Tight-VC-HC.

# A Algorithm Tight-VC-HC

In this section, we formally describe our tight approximation algorithm for VC-HC.

Let $\Pi = (V, E, c_v, m_v, d_e)$ denote the input instance of VC-HC. In the algorithm we start with a trivial initial parameter tuple $\Psi = (E, \mathbf{0}, \boldsymbol{c})$. As the algorithm iterates, the tuple $\Psi$ is refined gradually until we have a nice parameter tuple which makes Lemma 3 applicable. In this case we round up the solution given by LP($\Psi$) and stop. An overview of the algorithm is given in Figure 5. Below we describe the algorithm in details.

For a given parameter tuple $\Psi = (\mathcal{E}, \boldsymbol{\ell}, \boldsymbol{c'})$, the algorithm solves LP($\Psi$) for a basic optimal solution $p = (\boldsymbol{x}, \boldsymbol{h})$. Consider the set $I := \left\{v \colon v \in V,\ 0 < x_v < \frac{1}{f}\right\}$. For any $e \in \mathcal{E}$, let

$$T(e) := \left\{ v \ :\ v \in e \setminus I,\ 0 < h_{e,v} = x_v \right\}$$

denote the set of vertices in $V \setminus I$ which $e$ is supporting. If $\bigcup_{e\in\mathcal{E}} T(e)$ is empty, then no refinement is required and we proceed to the rounding stage. Otherwise, $\Psi$ is refined as follows:

- *Fold in supporting edges.* For each $e \in \mathcal{E}$ such that $T(e) \neq \emptyset$, we fold $e$ into $T(e)$ and modify $\Psi$ as follows. For each $v \in T(e)$ we subtract $d_e$ from $c'_v$ and set $\ell_v$ to be $x_v$. Afterwards we remove the edge $e$ from $\mathcal{E}$.

  Note that, in the folding process, we implicitly assign the edge $e$ to some particular $v \in T(e)$ and round down the assignment to zero for the remaining. Since $0 < h_{e,v} = x_v$ and

  $$d_e \cdot h_{e,v} \leq \sum_{e' \in \mathcal{E}[v]} d_{e'} \cdot h_{e',v} \leq c'_v \cdot x_v$$

  hold for any $v \in T(e)$, we have $d_e \leq c'_v$. Provided that vertex $v$ is to be rounded up later in the final stage, the validity of this process is thereby guaranteed.

If any refinement is made, this process repeats with the refined tuple of parameters. Otherwise we proceed to the rounding stage.

**Final rounding of the vertices.** Let $\tilde{p} = (\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{h}})$ denote the basic optimal solution computed before the algorithm enters the rounding stage and $\tilde{\Psi} = (\mathcal{E}, \boldsymbol{\ell}, \boldsymbol{c'})$ denote the corresponding parameter tuple. Let $\boldsymbol{h'}$ denote the rounded assignment function for the edges that are folded, i.e., for all $e \in E \setminus \mathcal{E}$, $h'_{e,v}$ is 1 for some particular $v \in T(e)$ and zero for the remaining.
The final rounding is done as follows. For any $v \in V$ and $e \in E[v]$, define

$$x^*_v := \left\lceil \tilde{x}_v \right\rceil \quad \text{and} \quad h^*_{e,v} := \begin{cases} \tilde{h}_{e,v}, & \text{if } e \in \mathcal{E}, \\ h'_{e,v}, & \text{otherwise.} \end{cases}$$

Then $(\boldsymbol{x}^*, \boldsymbol{h}^*)$ is output as the solution.

# B Proof of Theorem 2

To prove Theorem 2, we show that

- Algorithm Tight-VC-HC terminates in $O(|E|)$ rounds and outputs a feasible solution for $\mathrm{LP}(\Psi_0)$, where $\Psi_0 = (E, \mathbf{0}, \boldsymbol{c})$ is the initial tuple.
- The output solution $(\boldsymbol{x}^*, \boldsymbol{h}^*)$ is an $f$-approximation for VC-HC.

Consider the first statement to be proved. Since in each round of the execution, either the algorithm stops or at least one edge is folded. It follows that the algorithm must terminate in $O(|E|)$ rounds. Therefore it remains to prove the following:

  - After each round, the refined tuple is feasible, i.e., the corresponding LP has a non-empty feasible region. We prove this statement in Lemma 5.
  - The output solution $(\boldsymbol{x}^*, \boldsymbol{h}^*)$ is feasible for $\mathrm{LP}(\Psi_0)$. This is proved in Lemma 6.

Then we establish the approximation guarantee of $(\boldsymbol{x}^*, \boldsymbol{h}^*)$ in Section B.2.

**Notations and basic properties.** Let $k$ denote the number of rounds the algorithm iterates before entering the rounding stage. For $1 \le i \le k$, we use the following notations to denote the respective concepts we have in the $i^{th}$ iteration:

- $\Psi^{(i)} = \big(\mathcal{E}^{(i)}, \boldsymbol{\ell}^{(i)}, \boldsymbol{c}^{(i)}\big)$: The parameter tuple algorithm Tight-VC-HC maintains when it enters the $i^{th}$ rounds. Here we have $\Psi^{(1)} = \Psi_0$ as a dummy notation for the initial tuple.
- $p^{(i)} = \big(\boldsymbol{x}^{(i)}, \boldsymbol{h}^{(i)}\big)$: The basic optimal solution computed for $\mathrm{LP}(\Psi^{(i)})$.
- $I^{(i)}$: The set of vertices with small fractional values, i.e., $I^{(i)} := \{v \in V : 0 < x_v^{(i)} < \frac{1}{f}\}$.

The following proposition states a sandwich property for $\ell^{(i)}$ which we will be using later.

**Proposition 4.** For any $1 \le i < k$ and any $v \in V$, we have $\ell_v^{(i)} \le \ell_v^{(i+1)} \le x_v^{(i)}$. Furthermore, $\ell_v^{(i)} > 0$ implies that $\dfrac{1}{f} \le \ell_v^{(i)} \le 1$.

*Proof.* For the first statment, it suffices to see that, in the algorithm, $\ell_v^{(i+1)}$ equals either $\ell_v^{(i)}$ or $x_v^{(i)}$. In both cases it implies that $\ell_v^{(i)} \le \ell_v^{(i+1)} \le x_v^{(i)}$.

The second statement follows from the non-decreasing property of $\ell^{(i)}$ with respect to $i$ and the fact that $\ell_v^{(i)}$ is updated only when $v \notin I$ and $0 < x_v^{(i)} = h_{e,v}^{(i)}$ for some $e$, which implies that $\dfrac{1}{f} \le x_v^{(i)} \le 1$. □

## B.1 Feasibility of Tight-VC-HC

For any $1 \le i \le k$ and any edge subset $E' \subseteq E$, we define the extended assignment $\boldsymbol{h}^{(i)}|_{E'}$ of $\boldsymbol{h}^{(i)}$ with respect to the edge set $E'$ as follows. For any $e \in E'$ and any $v \in e$, let

$$\big(h^{(i)}|_{E'}\big)_{e,v} := \begin{cases} h_{e,v}^{(i)}, & \text{if } e \in \mathcal{E}^{(i)}, \\ 0, & \text{otherwise.} \end{cases}$$

The following lemma shows that, the feasible region of $\mathrm{LP}(\Psi^{(i+1)})$ is not empty for any $1 \le i < k$. Therefore the computation for $p^{(i)}$ results in a valid solution for all $1 \le i \le k$.

**Lemma 5.** $\left(\boldsymbol{x}^{(i)}, \boldsymbol{h}^{(i)}|_{\mathcal{E}^{(i+1)}}\right)$ is feasible for $\mathrm{LP}(\Psi^{(i+1)})$, for any $1 \leq i < k$.

*Proof.* We show that, if $p^{(i)}$ is feasible for $\mathrm{LP}(\Psi^{(i)})$, then the solution $\left(\boldsymbol{x}^{(i)}, \boldsymbol{h}^{(i)}|_{\mathcal{E}^{(i+1)}}\right)$ will be feasible for $\mathrm{LP}(\Psi^{(i+1)})$. Note that, the validity of the base case, $i = 1$, follows from the assumption that the input instance is feasible.
In the following, we show that $\left(\boldsymbol{x}^{(i)}, \boldsymbol{h}^{(i)}|_{\mathcal{E}^{(i+1)}}\right)$ does not violate the constraints in $\mathrm{LP}(\Psi^{(i+1)})$.

- Since $\mathcal{E}^{(i+1)} \subseteq \mathcal{E}^{(i)}$, it follows that $\left(h^{(i)}|_{\mathcal{E}^{(i+1)}}\right)_{e,v} = h^{(i)}_{e,v}$ for all $e \in \mathcal{E}^{(i+1)}$ and all $v \in e$. Therefore the constraints (1a) and (1d) hold directly for all $e \in \mathcal{E}^{(i+1)}$ and $v \in e$.

- Consider the constraint (1c). We have $x^{(i)}_v \leq m_v$ since $p^{(i)}$ is feasible for $\mathrm{LP}(\Psi^{(i)})$. By Proposition 4 we know that $\ell^{(i+1)}_v \leq x^{(i)}_v$. Hence the constraint (1c) holds for all $v \in V$.

- It remains to verify that the constraint (1b).

  For any $v \in V$, let $t^{(i)}(v)$ denote the set of edges in $\mathcal{E}^{(i)} \setminus \mathcal{E}^{(i+1)}$ that support $v$ in $p^{(i)}$. It follows that

  $$c^{(i)}_v - \sum_{e \in t^{(i)}(v)} d_e \;\leq\; c^{(i+1)}_v.$$

  Expanding constraint (1b) for $v$, we have

  $$\begin{aligned}
  \sum_{e \in \mathcal{E}^{(i+1)}[v]} d_e \cdot \left(h^{(i)}|_{\mathcal{E}^{(i+1)}}\right)_{e,v} &= \sum_{e \in \mathcal{E}^{(i)}[v]} d_e \cdot h^{(i)}_{e,v} \;-\; \sum_{e \in \mathcal{E}^{(i)}[v] \setminus \mathcal{E}^{(i+1)}[v]} d_e \cdot h^{(i)}_{e,v} \\
  &\leq \sum_{e \in \mathcal{E}^{(i)}[v]} d_e \cdot h^{(i)}_{e,v} \;-\; \sum_{e \in t^{(i)}(v)} d_e \cdot h^{(i)}_{e,v} \\
  &\leq \left(c^{(i)}_v - \sum_{e \in t^{(i)}(v)} d_e\right) \cdot x^{(i)}_v \\
  &\leq c^{(i+1)}_v \cdot x^{(i)}_v.
  \end{aligned}$$

  Therefore the constraint (1b) holds for all $v \in V$ and this lemma is proved. □

Let $\boldsymbol{h}'$ denote the assignment function defined in the algorithm for the edges that have been folded. The following lemma shows that the output solution $(\boldsymbol{x}^*, \boldsymbol{h}^*)$ is feasible for $\mathrm{LP}(\Psi_0)$.

**Lemma 6.** The solution $p = \left(\, \lceil \boldsymbol{x}^{(k)} \rceil \,,\, \boldsymbol{h}^{(k)}|_E + \boldsymbol{h}'|_E \,\right)$ is feasible for $\mathrm{LP}(\Psi_0)$.

*Proof.* We show that $p$ does not violate the constraints in $\mathrm{LP}(\Psi_0)$.

1. For the constraint (1a), it suffices to see that: (i) For any $e \in \mathcal{E}^{(k)}$, $(h'|_E)_{e,v} = 0$ for all $v \in e$. Therefore

   $$\sum_{v \in e} \left(h^{(k)}|_E + h'|_E\right)_{e,v} \;=\; \sum_{v \in e} h^{(k)}_{e,v} \;=\; 1.$$

   (ii) For any $e \in E \setminus \mathcal{E}^{(k)}$, we know that $(h^{(k)}|_E)_{e,v} = 0$ for all $v \in e$ and $\sum_{v \in e} h'_{e,v} = 1$ by the definition of $h'$. Hence constraint (1a) holds for all $e \in E$.

2. Consider the constraint (1b) for any $v \in V$. We consider two cases.

   (i) If $x^{(k)}_v = 0$, then we have $h^{(k)}_{e,v} = 0$ for all $e \in \mathcal{E}^{(k)}$.

Since $\ell_v^{(1)} = 0$, by Proposition 4 we have $\ell_v^{(i)} = 0$ for all $1 \le i \le k$. This means that $v$ is not supported by any edge in $p^{(i)}$ for all $1 \le i \le k$. Hence $(h'|_E)_{e,v} = 0$ for all $e \in E[v]$. Therefore,

$$\sum_{e \in E[v]} d_e \cdot \left(h^{(k)}|_E + h'|_E\right)_{e,v} \;=\; 0 \;=\; c_v \cdot \left\lceil x_v^{(k)} \right\rceil,$$

and the constraint (1b) holds for $v$.

(ii) Consider the case $x_v^{(k)} > 0$. We have

$$\begin{aligned}
\sum_{e \in E[v]} d_e \cdot \left(h^{(k)}|_E + h'|_E\right)_{e,v} \;&=\; \sum_{e \in \mathcal{E}^{(k)}[v]} d_e \cdot h_{e,v}^{(k)} \;+\; \sum_{e \in E[v] \setminus \mathcal{E}^{(k)}} d_e \cdot h'_{e,v} \\
&\le\; c_v^{(k)} \cdot x_v^{(k)} \;+\; \sum_{e \in E[v] \setminus \mathcal{E}^{(k)} :\, h'_{e,v} = 1} d_e \\
&\le\; \left( c_v^{(k)} + \sum_{e \in E[v] \setminus \mathcal{E}^{(k)} :\, h'_{e,v} = 1} d_e \right) \cdot \left\lceil x_v^{(k)} \right\rceil \;\le\; c_v \cdot \left\lceil x_v^{(k)} \right\rceil.
\end{aligned}$$

Therefore the constraint (1b) holds for $v$ as well.

3. Consider the constraint (1c) for any $v \in V$.

   By Proposition 4, we know that $\ell_v^{(i)}$ is non-decreasing on $i$. Therefore it follows that

   $$\left\lceil x_v^{(k)} \right\rceil \;\ge\; x_v^{(k)} \;\ge\; \ell_v^{(k)} \;\ge\; \ell_v^{(1)} \;=\; 0.$$

   On the other hand, since $m_v$ is integral and since $x_v^{(k)} \le m_v$, it follows that $\left\lceil x_v^{(k)} \right\rceil \le m_v$ as well. Hence constraint (1c) holds for all $v \in V$.

4. Consider the constraint (1d) and any $e \in E$, $v \in e$.

   If $e \in \mathcal{E}^{(k)}$, then we have $\left(h^{(k)}|_E + h'|_E\right)_{e,v} = h_{e,v}^{(k)} \le x_v^{(k)} \le \left\lceil x_v^{(k)} \right\rceil$ and the constraint (1d) holds naturally for $e, v$.

   If $e \notin \mathcal{E}^{(k)}$, then $\left(h^{(k)}|_E + h'|_E\right)_{e,v} = h'_{e,v}$. By the definition of $h'$, we know that $h'_{e,v} > 0$ implies that $0 < h_{e,v}^{(i)} = x_v^{(i)}$ and $\ell_v^{(i)} = x_v^{(i)}$ for some $1 \le i < k$. Therefore

   $$h'_{e,v} \;=\; 1 \;\le\; \left\lceil x_v^{(i)} \right\rceil \;=\; \left\lceil \ell_v^{(i)} \right\rceil \;\le\; \left\lceil \ell_v^{(k)} \right\rceil \;\le\; \left\lceil x_v^{(k)} \right\rceil.$$

   On the other hand, if $h'_{e,v} = 0$, then constraint (1d) holds trivially. In both cases, we know that it holds for $e$ and $v$. □

## B.2 Approximation Guarantee

In this section we show that $(\boldsymbol{x}^*, \boldsymbol{h}^*)$ gives an $f$-approximation. Since $x_v^* = \left\lceil x_v^{(k)} \right\rceil$ for all $v \in V$, we will prove in Lemma 8 that

$$\sum_{v \in V} \left\lceil x_v^{(k)} \right\rceil \;\le\; f \cdot \sum_{v \in V} x_v^{(1)}.$$

Consider the two sets $I^{(k)}$ and $D := \left\{ v :\, v \in V,\; 1 < x_v^{(k)} < m_v \right\}$. By their definitions and Proposition 4, we know that $I^{(k)}$ and $D$ are mutually disjoint, and they consist of only non-extremal vertices. Furthermore, $I^{(k)}$ is not supporting and $D$ is not supported.

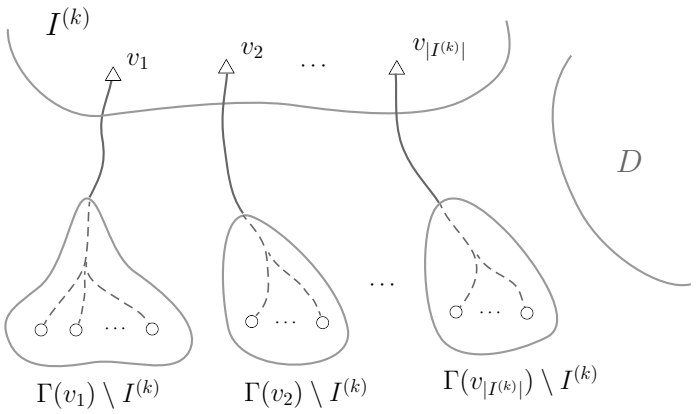


Figure 6: An illustration on the partition given by $\Gamma$.

Therefore, by Lemma 3, there exists a mapping

$$\Gamma\colon \;\; I^{(k)} \;\; \mapsto \;\; \Big(\mathcal{E}^{(k)}\Big)^{\text{actv}}_{h^{(k)}}\Big[I^{(k)}\Big] \;\setminus\; \Big(\mathcal{E}^{(k)}\Big)^{\text{actv}}_{h^{(k)}}\Big[D\Big] \tag{2}$$

such that: (i) $v \in \Gamma(v)$ for all $v \in I^{(k)}$, (ii) for any $u, v \in I^{(k)}$ with $u \neq v$, we have

$$\Big(\Gamma(u) \setminus I^{(k)}\Big) \cap \Big(\Gamma(v) \setminus I^{(k)}\Big) \;\; = \;\; \emptyset. \tag{3}$$

See also Figure 6 for an illustration. Note that, since we have $x_v^{(k)} < 1/f$ for all $v \in I^{(k)}$, it follows that

$$\sum_{v \in e \cap I^{(k)}} h_{e,v}^{(k)} < 1, \quad \text{and therefore} \quad e^{\text{actv}}_{h^{(k)}} \not\subseteq I^{(k)} \quad \text{for all } e \in \mathcal{E}^{(k)}.$$

Hence it follows that $\Gamma(v) \setminus I^{(k)} \neq \emptyset$ for all $v \in I^{(k)}$. In the following we will use the vertices in $\Gamma(v) \setminus I^{(k)}$ to help absorb the rounding cost of $v$.

For any $v \in I^{(k)}$, define $\pi(v) := \Big(\Gamma(v) \setminus I^{(k)}\Big) \cup \{v\}$. The following lemma bounds the rounding cost for the vertices included by an individual $\pi(v)$.

**Lemma 7.** Provided that $f \geq 2$, we have

$$\sum_{u \in \pi(v)} \Big\lceil x_u^{(k)} \Big\rceil \;\; \leq \;\; f \cdot \sum_{u \in \pi(v)} x_u^{(k)}, \quad \text{for any } v \in I^{(k)}.$$

*Proof.* From the Condition (2) above, we know that $\Gamma(v)$ contains no vertices in $D$. Since $\pi(v) \subseteq \Gamma(v)$, it follows that $\pi(v) \cap D = \emptyset$ as well. Depending on the elements of $\pi(v)$, we consider two cases.

- If there exists some $u \in \pi(v)$ with $x_u^{(k)} = m_u$, then we have

$$\Big\lceil x_v^{(k)} \Big\rceil + \Big\lceil x_u^{(k)} \Big\rceil \;\; = \;\; 1 + m_u \;\; \leq \;\; f \cdot \Big( x_v^{(k)} + x_u^{(k)} \Big),$$

  since $f \geq 2$ and $m_u \geq 1$. For the remaining vertices $u' \in \pi(v) \setminus \{u \cup v\}$, we know that $x_{u'}^{(k)} \geq 1/f$ since $u' \in \Gamma(v) \setminus I^{(k)}$ by definition of $\pi$. Hence $\Big\lceil x_{u'}^{(k)} \Big\rceil \leq f \cdot x_{u'}^{(k)}$, and the statement of this lemma holds.

- Consider the other case that $x_u^{(k)} \neq m_u$ for all $u \in \pi(v)$. Then we know that $x_u^{(k)} \leq 1$ for all $u \in \pi(v)$ and hence

$$\sum_{u \in \pi(v)} \left\lceil x_u^{(k)} \right\rceil \;=\; \Big|\pi(v)\Big|. \tag{4}$$

In the following we bound $|\pi(v)|$. Let $e \in \mathcal{E}^{(k)}[v]$ be an edge such that $\Gamma(v) = e^{\mathrm{actv}}_{h^{(k)}}$. Such an edge exists since $\Gamma(v)$ is the active subedge of certain edge in $\mathcal{E}^{(k)}[I^{(k)}]$ which also contains $v$. It follows that $\pi(v) \subseteq \Gamma(v) \subseteq e$.

We have

$$\sum_{u \in \pi(v)} x_u^{(k)} \;\geq\; \sum_{u \in \pi(v)} h_{e,u}^{(k)} \;=\; 1 - \sum_{u \in \Gamma(v) \setminus \pi(v)} h_{e,u}^{(k)} \tag{5}$$

where the first inequality follows from constraint (1d) and the second equality follows from constraint (1a) for $e$.

Since $\Gamma(v) \setminus \pi(v) \subseteq I^{(k)}$ by the definition of $\pi(v)$, it follows that

$$1 - \sum_{u \in \Gamma(v) \setminus \pi(v)} h_{e,u}^{(k)} \;>\; 1 - \frac{1}{f} \cdot \Big|\, \Gamma(v) \setminus \pi(v) \,\Big| \;\geq\; \frac{1}{f} \cdot \Big|\pi(v)\Big|, \tag{6}$$

where in the second inequality we use the fact that $|\Gamma(v)| \leq |e| \leq f$, which implies that $1 - \frac{1}{f} \cdot (\, |\Gamma(v)| - |\pi(v)| \,) \geq \frac{1}{f} \cdot |\pi(v)|$. Combining (4), (5), and (6), we obtain

$$\sum_{u \in \pi(v)} \left\lceil x_u^{(k)} \right\rceil \;=\; \Big|\pi(v)\Big| \;\leq\; f \cdot \sum_{u \in \pi(v)} x_u^{(k)}$$

as claimed. □

The following lemma establishes the approximation guarantee for the solution $(x^*, h^*)$.

**Lemma 8.** Provided that $f \geq 2$, we have

$$\sum_{u \in V} \left\lceil x_u^{(k)} \right\rceil \;\leq\; f \cdot \sum_{u \in V} x_u^{(1)}.$$

*Proof.* From the definition of $\pi$, we know that $\pi(u) \cap \pi(v) = \emptyset$ for any $u, v \in I^{(k)}$ with $u \neq v$. Furthermore, for any $u \in V \setminus \bigcup_{v \in I^{(k)}} \pi(v)$, we have $x_u^{(k)} \geq 1/f$ whenever $x_u^{(k)} > 0$.

Therefore, combining with Lemma 7, we have

$$\begin{aligned}
\sum_{u \in V} \left\lceil x_u^{(k)} \right\rceil \;&=\; \sum_{v \in I^{(k)}} \sum_{u \in \pi(v)} \left\lceil x_u^{(k)} \right\rceil \;+ \sum_{u \in V \setminus \bigcup_{v \in I^{(k)}} \pi(v)} \left\lceil x_u^{(k)} \right\rceil \\
&\leq\; \sum_{v \in I^{(k)}} f \cdot \sum_{u \in \pi(v)} x_u^{(k)} \;+\; f \cdot \sum_{u \in V \setminus \bigcup_{v \in I^{(k)}} \pi(v)} x_u^{(k)} \\
&\leq\; f \cdot \sum_{u \in V} x_u^{(k)}.
\end{aligned}$$

By Lemma 5, $\sum_{u \in V} x_u^{(i+1)} \leq \sum_{u \in V} x_u^{(i)}$ holds for all $1 \leq i < k$, since $x^{(i)}$ is feasible for $\mathrm{LP}(\Psi^{(i+1)})$ while $x^{(i+1)}$ is optimal for $\mathrm{LP}(\Psi^{(i+1)})$. Therefore

$$\sum_{u \in V} \left\lceil x_u^{(k)} \right\rceil \;\leq\; f \cdot \sum_{u \in V} x_u^{(k)} \;\leq\; f \cdot \sum_{u \in V} x_u^{(1)}$$

as claimed. □

# C Proof of Lemma 3

In this section we prove our main technical tool Lemma 3. We remark that, throughout this proof, we will assume the following prerequisite (as stated in the lemma):

1. $\Psi = (E, \boldsymbol{\ell}, \boldsymbol{c})$ denotes the considered parameter tuple,
2. $p = (\boldsymbol{x}, \boldsymbol{h})$ is an extreme point solution for $\mathbf{Q}(\Psi)$, and
3. $\mathcal{I}$ and $\mathcal{D}$ are two disjoint sets of non-extremal vertices (with reference to the point $p$) such that $\mathcal{I}$ is not supporting and $\mathcal{D}$ is not supported.

With the prerequisite above, the proof proceeds as follows. First in Section C.1 we identify a matrix $\tilde{M}$ from the constraint matrix of $\mathbf{Q}(\Psi)$ according to $\mathcal{I}$ and $\mathcal{D}$ together with the set of constraints they have involved in. From the matrix $\tilde{M}$ we show in Section C.2 that a mapping with nice and restricted behavior can be extracted for the final mapping $\Gamma$ to be defined in Section C.3.

## C.1 The constraint matrix of $\mathbf{Q}(\Psi)$ and the matrix $\tilde{M}$

Let $M$ denote the coefficient matrix for the constraints in $\mathbf{Q}(\Psi)$ and $M_{(\text{Eqs})}$ denote the submatrix formed by the constraints that hold with equality at the considered extreme point $p$. The following proposition states an equivalent form of a classical characterization on the extremality of $p$ which we will be using throughout the proof.

**Proposition 9.** *$M_{(Eqs)}$ has a rank equal to the number of variables in $\mathbf{Q}(\Psi)$.*

In the following, we first identify the sets of variables and constraints that are relative to $\mathcal{I}$ and $\mathcal{D}$. Then we show that, by properly simplifying the constraints, a matrix $\tilde{M}$ with a full column rank and nice structural property can be identified.

**Relative variables and constraints.** Consider two sets $X$ and $H$ of variables defined as

$$X := \{\ x_v :\ v \in \mathcal{I} \cup \mathcal{D}\ \} \quad \text{and} \quad H := \left\{\ h_{e,v} :\ e \in E[\mathcal{I}] \text{ such that } v \in e_h^{\text{actv}}\ \right\}.$$

Let $\text{Eqs}(X, H)$ denote the set of constraints in which the variables in $X \cup H$ have involved and which also hold with equality. We classify the equalities of $\text{Eqs}(X, H)$ by considering the four categories of constraints in $\mathbf{Q}(\Psi)$ as follows:

- For the constraint (1a), we use

$$\mathcal{C}^{(E)} := \left\{ e : \boxed{\textstyle\sum_{v \in e} h_{e,v} = 1} \in \text{Eqs}(X, H) \right\}$$

  to denote the the set of edges whose constraints are in $\text{Eqs}(X, H)$. Note that, it follows that $\mathcal{C}^{(E)} = E[\mathcal{I}]$, since for any $e \in E[\mathcal{I}]$, at least one variable of the constraint $\sum_{v \in e} h_{e,v} = 1$ is included in $H$ by definition.

- For the constraint (1b), we use

$$\mathcal{C}^{(V)} := \left\{ v : \boxed{\textstyle\sum_{e \in E[v]} d_e \cdot h_{e,v} \ \le\ c_v \cdot x_v} \in \text{Eqs}(X, H) \right\}$$

  to denote the set of vertices whose capacity constraints are in $\text{Eqs}(X, H)$.

- Since $\mathcal{I}$ and $\mathcal{D}$ are non-extremal, it follows that $\ell_v < x_v < m_v$ for all $x_v \in X$. Therefore the constraint (1c): $\ell_v \le x_v \le m_v$ does not appear in $\text{Eqs}(X, H)$ for all $v \in \mathcal{I} \cup \mathcal{D}$.

- For the constraint (1d), since $h_{e,v} > 0$ for all $h_{e,v} \in H$, the constraint $0 \leq h_{e,v}$ does not appear in $\mathrm{Eqs}(X, H)$ for all $e \in E$ and $v \in e$. On the other hand, we use

$$\mathcal{C}^{(E\times V)} := \left\{ \ (e,v) : \ x_v \in X, \ \ h_{e,v} \in H, \ \boxed{h_{e,v} \leq x_v} \in \mathrm{Eqs}(X,H) \ \right\}$$

  to denote the set of supporting constraints in $\mathrm{Eqs}(X, H)$. Since $\mathcal{I}$ is not supporting and $\mathcal{D}$ is not supported, we know that $(e, v) \in \mathcal{C}^{(E\times V)}$ implies that $v \in \mathcal{I}$.

Note that, from the classification above, it follows that $\mathrm{Eqs}(X, H) = \mathcal{C}^{(E)} \cup \mathcal{C}^{(V)} \cup \mathcal{C}^{(E\times V)}$. Let

$$H^* := H \setminus \left\{ h_{e,v} : \ (e,v) \in \mathcal{C}^{(E\times V)} \right\}$$

denote the set of variables in $H$ that does not correspond to a supporting constraint. For the simplicity of notations, we will also use $H^*$ to denote the set of pairs $(e, v)$ such that $h_{e,v}$ is contained in $H^*$, when there is no confusion in the context.

**The coefficient matrix $M_{(X,H)}$.** Let $M_{(X,H)}$ denote the submatrix of $M_{(\mathrm{Eqs})}$ that is formed by $X \cup H$ and $\mathrm{Eqs}(X, H)$.

$$M_{(\mathrm{Eqs})} = \begin{array}{c} \overbrace{\hphantom{M_{(X,H)}}}^{X,H} \\ \left(\begin{array}{c:c} \ddots & M_{(X,H)} \\ \hdashline \cdots & \mathbf{0} \end{array}\right) \end{array} \begin{array}{l} \left.\vphantom{\begin{array}{c}a\\a\end{array}}\right\} \mathrm{Eqs}(X,H) \\ \\ \end{array}$$

Provided the above classification for constraints in $\mathrm{Eqs}(X, H)$ and the definition of $H^*$, the matrix $M_{(X,H)}$ can be written as follows, where the submatrices $M^{(1)}_{\mathcal{C}^{(E\times V)}}$ and $M^{(2)}_{\mathcal{C}^{(E\times V)}}$ are described in Figure 7.

$$M_{(X,H)} = \begin{array}{c} \begin{array}{cccc} h_{e,v}, & h_{e,v}, & x_v, & x_v, \\ (e,v) \in \mathcal{C}^{(E\times V)} & (e,v) \in H^* & v \in \mathcal{I} & v \in \mathcal{D} \end{array} \\ \left(\begin{array}{c:c:c:c} \vdots & \vdots & \mathbf{0} & \mathbf{0} \\ \hdashline \mathbf{0} & & \mathbf{0} & M^{\mathcal{D}}_{\mathcal{C}^{(V)}} \\ \hdashline M^{(1)}_{\mathcal{C}^{(E\times V)}} & \mathbf{0} & M^{(2)}_{\mathcal{C}^{(E\times V)}} & \mathbf{0} \end{array}\right) \end{array} \begin{array}{l} \mathcal{C}^{(E)}: \ \sum h_{e,v} = 1 \\ \hdashline \mathcal{C}^{(V)}: \ \sum d_e \cdot h_{e,v} \leq c_v \cdot x_v \\ \hdashline \mathcal{C}^{(E\times V)}: \ \ h_{e,v} \leq x_v \end{array}$$

For the submatrix $M^{\mathcal{D}}_{\mathcal{C}^{(V)}}$, consider the column to which the variable $x_v$ corresponds for any $v \in \mathcal{D}$. Since $M_{(\mathrm{Eqs})}$ has full column rank and since $x_v$ does not appear in the constraints of $\mathcal{C}^{(E)}$ and $\mathcal{C}^{(E\times V)}$, it follows that $v$ must be in $\mathcal{C}^{(V)}$.[4] Therefore the matrix $M^{\mathcal{D}}_{\mathcal{C}^{(V)}}$ can be written as a diagonal matrix with coefficients $\left\{- c_v\right\}_{v\in\mathcal{D}}$ on its diagonals.

[4]Otherwise, the column $x_v$ would be a zero vector, a contradiction.

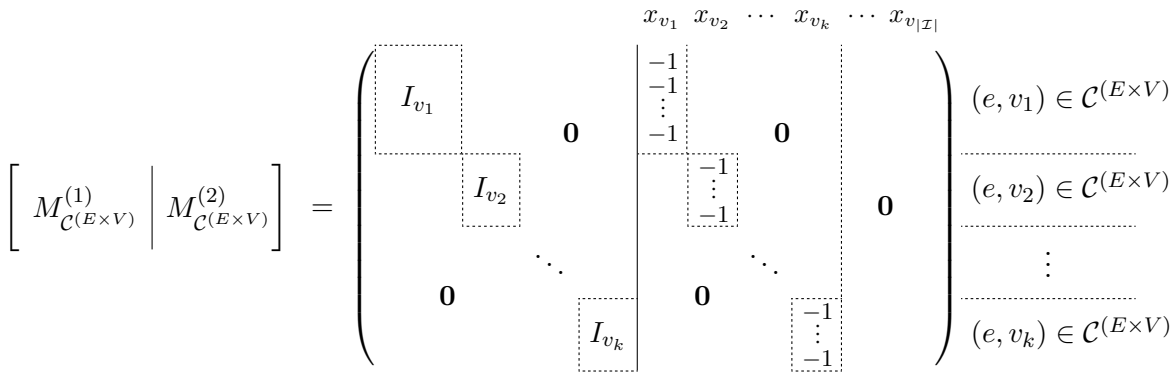


Figure 7: The matrices $M^{(1)}_{\mathcal{C}^{(E\times V)}}$ and $M^{(2)}_{\mathcal{C}^{(E\times V)}}$, where $I_{v_i}$ denotes the identity matrix with dimension $\left|\left\{e\colon (e, v_i) \in \mathcal{C}^{(E\times V)}\right\}\right|$.

**Reduced constraints $\tilde{\mathcal{C}}^{(E)}$ and $\tilde{\mathcal{C}}^{(V)}$, and the submatrix $\tilde{M}$.** Use the identity elements in the diagonal of $M^{(1)}_{\mathcal{C}^{(E\times V)}}$ as pivots and perform row reduction (Gaussian elimination on the rows) on $M_{(X,H)}$. Let the resulting matrix be $M'_{(X,H)}$. Note that, in this process, we literally replace the variable $h_{e,v}$ with $x_v$ in the constraints $\mathcal{C}^{(E)}$ and $\mathcal{C}^{(V)}$ for all $(e, v) \in \mathcal{C}^{(E\times V)}$. In particular, for each $e \in \mathcal{C}^{(E)}$, the constraint $\sum_{v\in e} h_{e,v} = 1$ is replaced by the constraint[5]

$$\sum_{\substack{v \text{ such that} \\ (e,v)\in H^*}} h_{e,v} \;+\; \sum_{\substack{v \text{ such that,} \\ (e,v)\in\mathcal{C}^{(E\times V)}}} x_v \;=\; 1. \tag{7}$$

Similarly, for each $v \in \mathcal{C}^{(V)}$, the original constraint is replaced by

$$\sum_{\substack{e\in E[v] \text{ such that} \\ (e,v)\notin\mathcal{C}^{(E\times V)}}} d_e \cdot h_{e,v} \;\le\; c'_v \cdot x_v, \quad \text{where } c'_v = c_v - \sum_{\substack{e \text{ such that,} \\ (e,v)\in\mathcal{C}^{(E\times V)}}} d_e. \tag{8}$$

We use $\tilde{\mathcal{C}}^{(E)}$ and $\tilde{\mathcal{C}}^{(V)}$ to denote the updated version of the constraints, i.e., (7) and (8), in $\mathcal{C}^{(E)}$ and $\mathcal{C}^{(V)}$, respectively.

Let $\tilde{M}$ denote the submatrix of $M'_{(X,H)}$ formed by the rows belonging to $\tilde{\mathcal{C}}^{(E)} \cup \tilde{\mathcal{C}}^{(V)}$ and the columns belonging to $X \cup H^*$. Lemma 10 follows from the extremality of $p$ and the fact that the constraints in $\mathcal{C}^{(E\times V)}$ are linearly independent.

**Lemma 10.** $\tilde{M}$ has full column rank.

*Proof.* By Proposition 9, the column vectors of the matrix $M_{(Eqs)}$ are linearly independent. Therefore, the column vectors in the submatrix $M_{(X,H)}$ are also linearly independent as they do not involve in constraints other than those in $\mathrm{Eqs}(X, H)$.

Since $M_{(X,H)}$ contains an identity matrix, i.e., $M^{(1)}_{\mathcal{C}^{(E\times V)}}$, by Gaussian elimination we know that, taking out the rows and the columns that belong to $M^{(1)}_{\mathcal{C}^{(E\times V)}}$ from $M_{(X,H)}$ would have decreased its rank by exactly the dimension of $M^{(1)}_{\mathcal{C}^{(E\times V)}}$, i.e., $\left|\,\mathcal{C}^{(E\times V)}\,\right|$.

[5]Note that, in Equation (7) we also drop out $h_{e,v}$ for all $e \in \mathcal{C}^{(E)}$ and $v \in e \setminus e_h^{\mathrm{actv}}$ since by definition they are zero at the considered extreme point $p$.

$$M'_{(X,H)}\colon \left(\begin{array}{c|c} \mathbf{0} & \tilde{M} \\ \hline M^{(1)}_{\mathcal{C}^{(E\times V)}} & \cdots \end{array}\right) \implies \left(\begin{array}{c|c} \mathbf{0} & \tilde{M} \\ \hline M^{(1)}_{\mathcal{C}^{(E\times V)}} & \mathbf{0} \end{array}\right)$$

Therefore $\tilde{M}$ has a full column rank as well. □

## C.2 Mapping obtained from extremality of $p$

In the following we consider the matrix $\tilde{M}$. Since $\tilde{M}$ has full column rank, for each column $j$ of $\tilde{M}$ there exists a distinct row $i$ such that $\tilde{M}(i,j)$ is non-zero[6], i.e., the existence of distinct pivot for each column. Let

$$\sigma\colon X \cup H^* \mapsto \tilde{\mathcal{C}}^{(E)} \cup \tilde{\mathcal{C}}^{(V)}$$

denote one of such mappings. Note that $\sigma$ is injective, i.e., for all $r, s \in X \cup H^*$, $r \neq s$ implies that $\sigma(r) \neq \sigma(s)$. For each $x_v \in X$, define a mapping $\pi\colon X \mapsto X \cup H^*$ as follows.

$$\pi(x_v) := \begin{cases} h_{e,v}, & \text{for } \textit{any } e \text{ such that } (e,v) \in H^*, \quad \text{if } v \in \mathcal{I} \text{ and } \sigma(x_v) = v, \\ x_v, & \text{otherwise.} \end{cases}$$

The following lemma shows that $\pi$ is a well-defined function of mapping.

**Lemma 11.** For any $v \in \mathcal{I}$ with $\sigma(x_v) = v$, there exists $e$ such that $(e,v) \in H^*$.

*Proof.* Let $v \in \mathcal{I}$ be a vertex with $\sigma(x_v) = v$. Since the image of $\sigma$ is $\tilde{\mathcal{C}}^{(E)} \cup \tilde{\mathcal{C}}^{(V)}$, from the assumption that $\sigma(x_v) = v$, we know that $v \in \tilde{\mathcal{C}}^{(V)}$.

Since $v \in \mathcal{I}$, we know that the variables $h_{e,v}$ for all $e \in E[v]$ with $h_{e,v} > 0$ have been included in $H$ by its definition. Therefore the Constraint (8) for $v$ can be written as

$$\sum_{\substack{e \text{ such that} \\ (e,v) \in H \setminus \mathcal{C}^{(E\times V)}}} d_e \cdot h_{e,v} \;=\; c'_v \cdot x_v. \tag{9}$$

If $(e,v) \in \mathcal{C}^{(E\times V)}$ for all $e$ such that $(e,v) \in H$, then the left hand side of (9) vanishes and we have $c'_v \cdot x_v = 0$. Since $v \in \mathcal{I}$, $v$ is non-extremal and it follows that $x_v > \ell_v \geq 0$. Therefore $c'_v$ must be zero. This means that, the row to which $v$ corresponds in $\tilde{M}$ is a zero vector, which renders it impossible to be mapped to by $\sigma$ since it has no pivots for the column $x_v$, a contradiction. Hence, there must exist $e$ such that $(e,v) \in H$ and $(e,v) \notin \mathcal{C}^{(E\times V)}$, which in turn implies that $(e,v) \in H^*$. □

Consider the mapping $(\sigma \circ \pi)\colon X \mapsto \tilde{\mathcal{C}}^{(E)} \cup \tilde{\mathcal{C}}^{(V)}$. Since both $\sigma$ and $\pi$ are injective, $(\sigma \circ \pi)$ is also injective. Lemma 12 provides an exact classification on the image of $(\sigma \circ \pi)$.

**Lemma 12.** $(\sigma \circ \pi)(x_v) \in \tilde{\mathcal{C}}^{(E)}$ for all $v \in \mathcal{I}$. In contrast, $(\sigma \circ \pi)(x_v) = v$ for all $v \in \mathcal{D}$.

*Proof.* First we show that $(\sigma \circ \pi)(x_v) \in \tilde{\mathcal{C}}^{(E)}$ if $v \in \mathcal{I}$. Depending on whether or not $\sigma(x_v) = v$, we consider two cases. If $\sigma(x_v) = v$, then it suffices to show that $\sigma(h_{e,v}) = e$ for all $e$ such that $(e,v) \in H^*$. Indeed, for any $e$ such that $(e,v) \in H^*$, provided that $v \in \tilde{\mathcal{C}}^{(V)}$ (since it is mapped by $\sigma$), the variable $h_{e,v}$ appears in exactly two constraints in $\tilde{\mathcal{C}}^{(E)} \cup \tilde{\mathcal{C}}^{(V)}$, i.e., the constraint $e$ and the constraint $v$, respectively. Since the constraint $v$ is already occupied by $x_v$ (in the mapping $\sigma$), it leaves $e$ the only choice for $h_{e,v}$ to occupy. See also Figure 8(a) for an illustration.

[6] If not, the column reduction (Gaussian elimination on the columns) would have led to a rank less than $|X \cup H^*|$, a contradiction to the fact that $\tilde{M}$ has a full column rank.

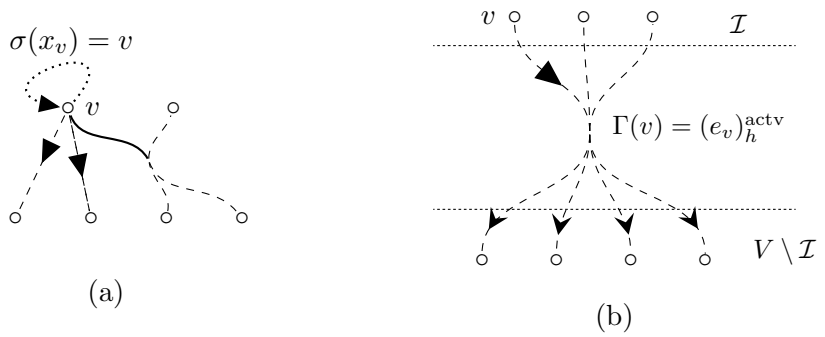


Figure 8: The restricted behavior of the mapping $\sigma$. The arrow head shows the target to which $\sigma$ is mapping. (a) If the variable $x_v$ is mapped to $v$ itself, then there exists non-empty outward mapping to its incident edges. (b) When an edge constraint is occupied (mapped) by the variable $\pi(x_v)$, the remaining variables in the edge constraint have to map outward.

If $\sigma(x_v) \neq v$, then $(\sigma \circ \pi)(x_v) = \sigma(x_v)$ by definition. Since the variable $x_v$ involves in exactly one constraint in $\tilde{\mathcal{C}}^{(V)}$, i.e., the constraint $v$ itself, $\sigma(x_v) \neq v$ implies that $\sigma(x_v)$ must be in $\tilde{\mathcal{C}}^{(E)}$. This proves the first part of the statement.

To see that $(\sigma \circ \pi)(x_v) = v$ for all $v \in \mathcal{D}$, it suffices to see that in the matrix $M_{(X,H)}$, and hence the matrix $\tilde{M}$, the column belonging to $x_v$ has one and only one non-zero element that is located in the row belonging to the constraint $v \in \tilde{\mathcal{C}}^{(V)}$, i.e., in the diagonal matrix $M^{\mathcal{D}}_{\mathcal{C}^{(V)}}$. Therefore we have $(\sigma \circ \pi)(x_v) = \sigma(x_v) = v$. □

Lemma 12 shows that the mapping $(\sigma \circ \pi)$ gives a distinct edge for each $v \in \mathcal{I}$. The following lemma provides properties for the "active ends" of these particular edges. Recall that $e_h^{\text{actv}}$ denotes the set of the active ends of the edge $e$.

**Lemma 13** (*Restricted Behavior of the Mappings*)**.** Consider the edge $e_v := (\sigma \circ \pi)(x_v)$ defined for any $v \in \mathcal{I}$. We have

1. $v \in \big(e_v\big)_h^{\text{actv}}$ for all $v \in \mathcal{I}$.
2. $\sigma(h_{e_v,u}) = u$ for any $u \in \big(e_v\big)_h^{\text{actv}} \setminus \mathcal{I}$.

*Proof.* Consider a vertex $v \in \mathcal{I}$ and the edge $e_v := (\sigma \circ \pi)(x_v)$. From the Constraint (7) for $e_v$ we know that $e_v$ is mapped either by the variable $h_{e_v,v}$ with $(e_v, v) \in H^*$ or the variable $x_v$, which in turn means that $(e_v, v) \in \mathcal{C}^{(E \times V)}$. Therefore $(e_v, v) \in H$ in both cases. It follows that $h_{e_v,v} > 0$ and $v \in \big(e_v\big)_h^{\text{actv}}$.

Below we prove the second part of this lemma. Since the constraint to which $e_v$ corresponds is already occupied by $\pi(x_v)$ (in the mapping $\sigma$), it follows that, the remaining variables of $X \cup H$ that appear in constraint $e_v$ must be mapped to constraints other than $e_v$. In particular, this statement holds for any $u \in \big(e_v\big)_h^{\text{actv}} \setminus \mathcal{I}$. If such a vertex $u$ exists, the variable $h_{e_v,u}$ would have been contained in $H$ since $\tilde{\mathcal{C}}^{(E)} = E[\mathcal{I}]$ and $u \in \big(e_v\big)_h^{\text{actv}}$. Since $\mathcal{I}$ is not supporting, it follows that $h_{e_v,u}$ involves in at most two constraints, i.e., $e_v$ and $u$. Since $e_v$ is already mapped by $\pi(x_v)$, it follows that $u$ must be included in $\tilde{\mathcal{C}}^{(V)}$ for $h_{e_v,u}$ to be mapped to. This shows that $\sigma(h_{e_v,u}) = u$. Figure 8(b) illustrates this argument. □

### C.3 The mapping $\Gamma$

For each $v \in \mathcal{I}$, define $\Gamma(v) := \big((\sigma \circ \pi)(x_v)\big)_h^{\text{actv}}$. Below we show that $\Gamma$ certifies the statement of Lemma 3.

1. (Reflexive)

   This follows directly from Lemma 13 and the definition of $\Gamma$.

2. (The image)

   Consider a vertex $v \in \mathcal{I}$. Since $v \in \big((\sigma \circ \pi)(x_v)\big)_h^{\text{actv}} = \Gamma(v)$ by the reflexive property, it follows that

   $$\Gamma(v) \in E_h^{\text{actv}}[v] \subseteq E_h^{\text{actv}}[\mathcal{I}]. \tag{10}$$

   For any $u \in \mathcal{D}$, from Lemma 12 we know that $\sigma(x_u) = u$. Therefore, by Lemma 13, $u$ cannot belong to $\Gamma(v)$ since $u \notin \mathcal{I}$ and since the constraint $u$ is already occupied by $x_u$ in the mapping $\sigma$. This implies that $\Gamma(v) \notin E_h^{\text{actv}}[u]$. Since this holds for all $u \in \mathcal{D}$, we have $\Gamma(v) \notin E_h^{\text{actv}}[\mathcal{D}]$. Combined with (10), it follows that

   $$\Gamma(v) \in E_h^{\text{actv}}[\mathcal{I}] \setminus E_h^{\text{actv}}[\mathcal{D}].$$

3. (Closed under intersection.)

   Consider any $u, v \in \mathcal{I}$ such that $u \neq v$. For any $w \in \Gamma(u) \cap \Gamma(v)$, if $w \notin \mathcal{I}$, then it follows from Lemma 13 that $\sigma(h_{(\sigma\circ\pi)(x_u),w}) = w$ and $\sigma(h_{(\sigma\circ\pi)(x_v),w}) = w$, a contradiction to the injective property of $\sigma$. Therefore $\Gamma(u) \cap \Gamma(v) \subseteq \mathcal{I}$.

This proves Lemma 3.